



Trigger & DAQ Development



Wesley H. Smith

U. Wisconsin – Madison


Instrumentation Frontier Meeting

U. Minnesota, July 31, 2013

with input from Su Dong, Gunther Haller, Mike Huffer, Ted Liu

Outline:

- **Trigger & DAQ Challenges**
- **Strategies for Experiments**
- **Tools: FPGAs, AM, xTCA, Transceivers, GPU, PCIe**
- **Directions for R&D**



Trigger & DAQ Challenges



Energy Frontier


- LHC: ATLAS & CMS
 - Highest data volumes & processing rates

Intensity Frontier

- LBNE: up to 0.8 Tbytes/s
- LHCb: 9 Tbytes/s
- Belle II, Mu2e...

Cosmic Frontier

- LSST: 3 Gbytes/s
- DArkSide: 100's of Mbytes/s, Gbytes/s for calibration
- CTA, CDMS, LZ...



LHC Experiment Scenarios



ALICE (post-LS2): Triggerless

- Readout 50 kHz Pb-Pb (*i.e.* $L = 6 \times 10^{27} \text{ cm}^{-1} \text{ s}^{-1}$), with minimum bias (pipeline) readout (max readout at present $\sim 500 \text{ Hz}$)

ATLAS (post-LS3): Triggered

- Divide L1 Trigger into L0, L1 of latency 5, 20 μsec , rate $< 500, 200 \text{ kHz}$, HLT output rate of 5 - 10 kHz
- L0 uses Calo & Muon Triggers, generates track trigger seeds
- L1 uses Track Trigger & more muon detectors & more fine-grained calorimeter trigger information.

CMS (post LS3): Triggered

- Considering L1 Trigger latency, rate: 10 – 20 μsec , 0.5 – 1 MHz
- L1 uses Track Trigger, finer granularity μ & calo. Triggers
- HLT output rate of 10 kHz

LHCb (post LS2): Triggerless

- Execute whole trigger on CPU farm \Rightarrow provide $\sim 40 \text{ MHz}$ detector readout



ATLAS & CMS Triggered vs. Triggerless Architectures




1 MHz (Triggered):

- **Network:**
 - 1 MHz with 10 MB: aggregate 80 Tbps
 - Links: Event Builder-cDAQ:
 - ~10,000 links of 10 Gbps or 1000 of 100 Gbps
 - Switch: almost possible today, for 2022 no problem
- **HLT computing:**
 - General purpose computing: $10(\text{rate}) \times 2(\text{PU}) \times 200\text{kHS6}$
 - Factor 20 wrt today
 - Maybe for ~same costs
 - Specialized computing (GPU or else)
 - Possible

40 MHz (Triggerless):

- **Network:**
 - 40 MHz with 10 MB: aggregate ~3,000 Tbps
 - Event Builder Links:
 - ~10,000 links of 100 Gbps
 - Switch: has to grow by factor ~60 in 10 years, not excluded but not likely
 - Readout Cables: Copper Tracker!
- **HLT computing:**
 - General purpose computing: $400(\text{rate}) \times 2(\text{PU}) \times 200\text{kHS6}$
 - Factor 800 wrt today
 - Looks impossible with realistic budget
 - Specialized computing (GPU or else)"
 - Could possibly provide this ...



HL-LHC Track Trigger Architectures:



“Push” path:


- L1 tracking trigger data combined with calorimeter & muon trigger data regionally with finer granularity than presently employed.
- After regional correlation stage, physics objects made from tracking, calorimeter & muon regional trigger data transmitted to Global Trigger.

“Pull” path:

- L1 calorimeter & muon triggers produce a “Level-0” or L0 “pre-trigger” with request for regional tracking info at ~1 MHz.
- Tracker sends out info. for these regions only & this data is combined in L1 correlation logic, resulting in L1A combining track, muon & cal. info..
- Only on-detector tracking trigger logic in specific region would see L0

“Afterburner” path:

- L1 Track trigger info, along with rest of information provided to L1 is used at first stage of HLT. Provides track information to HLT algorithms very quickly without having to unpack & process large volume of tracker information through CPU-intensive algorithms. Helps limit need for significant additional processor power in HLT computer farm.



HL-LHC L1 Trig. Latency, Rate



Latency: Provides option of simpler tracking trigger

- Timing is very tight for tracking trigger
 - Including processing & use of track trigger information
- “Pull” option
 - May want to keep advantages of “push” design anyway
- Makes design of tracking trigger easier
 - Relaxed constraints: reduces power, transmission bandwidth...

Latency: Provides option of pixel tracking trigger

- Pixel trigger requires “pull” architecture
- Required for b-tags in L1 Trigger
 - Along with 0.5-1 MHz L1 bandwidth

Rate: Reduces Thresholds for physics signals

- Can set thresholds comparable to present ones when coupled with tracking triggers

Rate: Needed for Hadronic Triggers

- Track Trigger helps leptonic triggers
- Less of an impact on hadronic triggers
 - Vertex for jets

Rate: Needed for b-tags

- Pixel trigger may not reduce rate sufficiently



HL-LHC HLT Output Rate



Processing 0.5-1 MHz Input


- DAQ hardware & HLT processing compatible with Moore's Law scaling until 2023 & estimated x3 longer reconstruction time, event size.
 - CMS predicts CPU time/event = 600 ms at PU=125 (200 now)
- Use of L1 Track Trigger information as input allows immediate, fast use of tracking information.
- Possibility to share resources with Tier-0 (Cloud computing)
 - Goes both ways
- If we need more CPU, we can bring more online rapidly if we can afford it (have already done this)

5-10 kHz Output Rate

- 1 MHz L1 Accept Rate \rightarrow 10 kHz HLT output rate keeps same reduction of L1 rate (x100) as present HLT design (100 kHz \rightarrow 1 kHz)
- Output to Computing
 - Compatible with Moore's Law scaling until 2023 & estimated X3 longer reconstruction time, event size



LHCb Upgrade Trigger & DAQ



1 – 40 MHz
All detectors information



20 kHz

Execute whole trigger on CPU farm

→ Provide ~40 MHz detector readout

- Cannot satisfy present 1 MHz requirement w/o deeply cutting into efficiency for hadronic final states
 - worst state is $\phi\phi$, but all hadronic modes are affected
 - Can ameliorate this by reading out detector & then finding vertices

Upgrade Trigger & DAQ

- flexible software trigger with up to 40 MHz input rate and 20 kHz output rate
- run at ~ 5-10 times nominal LHCb luminosity $\rightarrow L \sim 1-2 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$
- big gain in signal efficiency (up to x7 for hadron modes)
- upgrade electronics & DAQ architecture
- collect $\geq 5/\text{fb}$ per year and $\sim 50/\text{fb}$ in 10 years



ALICE Upgrade



Run at high rates, 50 kHz Pb-Pb (*i.e.* $L = 6 \times 10^{27} \text{ cm}^{-1} \text{ s}^{-1}$), with minimum bias (pipeline) readout (max readout with present ALICE set-up $\sim 500 \text{ Hz}$)

- Factor 100 increase in recorded luminosity
- Improve vertexing and tracking at low p_t

Pb-Pb run complemented by p-Pb & pp running

Entails building High-rate upgrade for readout of TPC, TRD, TOF, CALs, Muons, DAQ/HLT

Two HLT scenarios for the upgrade:


- Partial event reconstruction (clustering and tracking):
Factor of $\sim 20 \rightarrow$ Rate to tape: 20 kHz
 - clusters (associated with tracks) information recorded on tape
- Full event reconstruction:
additional reduction factor $\sim 3 \rightarrow$ Rate to tape $> 50 \text{ kHz}$
 - track parameters recorded on tape



SuperKEKB / Belle2 (2016)





Lumi: 8×10^{35}
Beam crossing: 4ns
L1: ~ 30 kHz
Logging rate: 6 kHz
Event size: 300 KB



New in Belle II

Global DAQ Design





Triggerless:


- Front end chips are installed in LAr to minimize the capacitance and noise
- On chip digitization to convert to digital signals inside detector cryostat
- Multiplexing to high speed serial link, to reduce cable plants, minimize outgassing, make possible the scalability to larger detector volumes
 - Balance with inaccessibility, programmability (ASICs in cold volume – FPGAs may not be reliable in this environment)

Also: DArkside:

- Transmits analog data from the cold, digitizes data in the warm, and uses RCE to process triggerless data



Mu2e DAQ - Triggerless



Streaming DAQ system

- 30 GBytes/sec bandwidth
- 30 TFLOPS processing
- simple architecture
- commodity hardware
- common software



LSST Block diagram & data-paths of the DAQ




Systems:

• Science & Wavefront

- Acquire, (re)format, cross-talk correct & deliver to multiple clients
- 4000 receive, process, store & transmit at > 3.2 GBytes/Sec

• Guider

- Delivers windowed data to multiple clients
- Telescope control systems






CTA Camera Trigger & DAQ

L0 trigger signals from target ASICs in camera modules are presented to Trigger FPGAs which use simple combinatorial logic to identify patterns

Trigger FPGA latches L0 hit pattern & local time & delivers to L2 camera trigger.





Tools for Triggers: FPGAs





Logic Cells

- 28 nm: > 2X gains over 40 nm →

On-Chip High Speed Serial Links:

- Connect to new compact high density optical connectors (SNAP-12...)






Tool for Tracking Triggers: Associative Memories

Pattern Recognition Associative Memory (PRAM)

- Based on CAM cells to match and majority logic to associate hits in different detector layers to a set of pre-determined hit patterns
- highly flexible/configurable, much less demand on detector design
- Pattern recognition finishes soon after hits arrive
- Potential candidate for L1 pattern recognition
- However: Latency
- Challenges:

- Increase pattern density by 2 orders of magnitude
- Increase speed x 3
- Same Power
- Use 3D architecture: Vertically Integrated Pattern Recognition AM - VIPRAM





Tools for Trigger/DAQ: ATCA






- **Advanced Telecommunications Computing Architecture ATCA**
- **Example: ATLAS Upgrade Calorimeter Trigger Topological Processor Card**

- 12-chan. ribbon fiber optic modules
- Backpl. opt. ribbon fiber connector

- **Example: μ TCA derived from AMC std. used by CMS HCAL, Trig.**

- **Advanced Mezzanine Card**
- **Up to 12 AMC slots**
 - *Processing modules*
- **6 standard 10Gb/s point-to-point links slot to hub slots (more available)**
- **Redundant power, controls, clocks**
- **Each AMC can have in principle (20) 10 Gb/sec ports**
- **Backplane customization is routine & inexpensive**





DAQ Tools: RCE System

Integrated hardware + software entity where generic core firmware & software infrastructure are common & provided.

ATCA infrastructure used


Xilinx ZYNQ series with ARM processors that can run either RTEMS or LINUX.

Has three principal components:

- Programmable FPGA Fabric
- Programmable Cluster-Element (CE).
- *Plugins*

Currently being used in:

- ATLAS CSC (proposed: Small Wheel), DArkside, Heavy Photon Search, LBNE, LSST, LCLS, nEXO...





Tools cont'd: CPU, GPU, PCIe



CPU Gains for High Level Triggers: Moore's Law GPU Enhancement of HLT →




- GPU performance tracks Moore's Law, since GPU architecture is scalable:

- Large Increase in memory bandwidth x10 in Gbytes/s
- Power efficient x3 with latest GPU card
- Well suited to tracking, fitting algorithms




Peak Double Precision FP



Enhancement of detector to DAQ readout:

- PCI Express Gen3 Cards now available
- Up to 56Gb/s InfiniBand or 40 Gigabit Ethernet per port



R&D Topics: Trigger (mostly ATLAS & CMS)



Increase of rate from Level-0 to HLT to read out

- Absolute rate & balance between levels

L1 complexity vs. HLT input rates

- Study the trade-offs

L1 Trigger Latency

- How much is needed & consequences on electronics

L1 Track Triggers

- Associative Memories
- Study techniques: sharpen p_T threshold, e- & μ - ID, Isolation, primary vertex for jets, multi-object triggers, possibility of pixel b-tag.
- Interplay with tracker design

Improvements to L1 Calo. & Muon Triggers

- Processing of much finer-grain, higher bandwidth information

Impact of higher bandwidth links & denser optical interconnects


New packaging & interconnect technologies

- ATCA, μ TCA, RCE

Use of FPGAs in L1 Trigger

Impact of detector timing improvements (~ 100 ps)

- e.g. crystal calorimeters (CMS: PbWO₃ has ~ 150 ps, LYSO < 100 ps)



R&D Topics: HLT & DAQ



New packaging & interconnect technologies

- ATCA, μ TCA, RCE

Event building architectures

Future of Server PC architecture

Network Switches

Clock & Control Networks

HLT on the Cloud

- e.g. share resources between HLT & Tier-0

HLT Specialized Track Processing

- e.g. GPU
- depends on resources available: cpu but also link speed

Simulation of HLT

- More sophisticated algorithms, increased occupancy

Use of New Processors in HLT

- ARM, Nvidia Tesla (GPU), Xeon Phi...
 - Just a list of what we can use in the next 2 years!
 - Eventually: heterogeneous mixtures of cores: general & specialized?
- Applies also to computing & software topics

Merging of HLT & offline software development